

The Fox Project: Advanced Development of Systems Software

R&D Status Report
October 31 to December 31, 1998

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

19990126 005

DISTRIBUTION STATEMENT A

**Approved for public release;
Distribution Unlimited**

This research is sponsored by the Defense Advanced Research Projects Agency, DoD, through ARPA Order 8313, and monitored by ESD/AVS under contract F19628-95-C-0050. Views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the United States Government.

DTIC QUALITY INSPECTED 1

The long-term objectives of the Carnegie Mellon Fox Project are to improve the design and construction of systems software and to further the development of advanced programming language technology. We use principles and techniques from the mathematical foundations of programming languages, including semantics, type theory, and logic, to design and implement systems software, including operating systems, network protocols, and distributed systems. Much of the implementation work is conducted in the Standard ML (SML) language, a modern functional programming language that provides polymorphism, first-class functions, exception handling, garbage collection, a parameterized module system, static typing, and a formal semantics. This Project involves several faculty members and spans a wide range of research areas, from (1) advanced compiler development to (2) language design to (3) software system safety infrastructure.

1 Research Progress

We report on the research accomplishments during the fourth calendar quarter of 1998, and the research objectives for the first quarter of 1999.

Accomplishments (October-December):

- Completed the first self-compile of TILT with the internal language type checking enabled. This demonstrates the use of typed intermediate languages to build a certifying compiler for a full-scale language.
- Some initial implementation work on parallelizing the TILT ML compiler has been accomplished. In particular, the language has been extended with a parallel-let construct, the compiler can reduce these parallel-lets to more primitive constructs like fork and spawn, and the runtime system has been augmented with a scheduler.
- Implemented prototype cost-based vcgen for PCC resource bounds research.
- Perry Cheng successfully proposed a thesis topic: Scalable Parallel Garbage Collection on Symmetric Multiprocessors.
- Significant progress has been made on the MLBox compiler implementation. Specifically, having abandoned home-grown code-generator and register allocator, we have tied compiler to the MLRISC system. We have successfully used MLRISC to generate static code. Currently using MLRISC to compile snippets of code that will be specialized and linked dynamically at run time.

- Completed implementation of the core typechecker for KML, a dialect of ML designed in part to experiment with language features in consideration for ML2000. In particular, the implementation illustrates in a practical setting the strengths and weaknesses of local type inference, a technology on which the current designs for ML2000 fundamentally depend.
- Devised type systems for intermediate languages capable of expressing low-level intensional type analysis, direct memory management, and efficient object representations, and devised formal compiler transformations for generated code in these intermediate languages. These languages make it possible to express optimizations in type-safe languages that previously were not permitted by any safe type system, bringing us closer to the goal of a fully typed compiler capable of state-of-the-art optimizations.

Objectives (January-March):

- Address serious theoretical and performance problems with the implementation of the type checkers for the TILT intermediate languages.
- Continue PCC resource bounds research.
- Work on the port of PCC to the x86 architecture.
- Complete a working MLBox compiler.
- Continue to work on the KML compiler implementation, moving toward linking it with the TILT compiler.
- Begin implementation of the new intermediate type systems and extend them with support for additional optimizations.

2 Noteworthy Publications

- *What is a Recursive Module?* by Karl Crary, Robert Harper and Sidd Puri. Accepted for publication in PLDI '99. Also published as technical report CMU-CS-FOX-98-03.
- *Relational Interpretations of Recursive Types in an Operational Setting* by Lars Birkedal and Robert Harper. Accepted for publication in Theoretical Computer Science B.

- *Transparent and Opaque Interpretations of Datatypes* by Karl Crary, Robert Harper, Perry Cheng, Leaf Petersen and Christopher Stone. Submitted for publication to the Journal of Functional Programming. Also published as technical report CMU-CS-FOX-98-04.
- *Parametricity and Variants of Girard's J Operator* by Robert Harper and John C. Mitchell. Submitted for publication to Information Processing Letters.
- *Privacy via Subsumption* by Jon G. Riecke and Christopher A. Stone. Accepted for publication in Theory and Practice of Object Systems.
- *Scalably Parallel Garbage Collection* by Guy Blelloch and Perry Cheng. Accepted for publication in PLDI '99.
- *Typed Memory Management in a Calculus of Capabilities* by Karl Crary, David Walker and Greg Morrisett. Accepted for publication in POPL '99.
- *Stack-Based Typed Assembly Language* by Greg Morrisett, Karl Crary, Neal Glew and David Walker. Submitted for publication to the Journal of Functional Programming. Also published as technical report CMU-CS-FOX-98-05.
- *Dependent Types in Practical Programming* by Frank Pfenning and Hongwei Xi. Accepted for publication in POPL '99.

3 Capital Equipment Purchases

- 1 Pentium II 450MHz Desktop Workstation; 1 Seagate 12-24GB DDS3 DAT Drive, \$3,475.00
- 4 Cisco 762 Ethernet/ISDN/NT1/IP Router – No. Am; 4 Cisco 760/770 Series Remote Office Feature Pack, \$2,417.20
- 1 IBM Thinkpad 600 Model 300, \$3,266.00
- 1 Hitachi 21" XGA Monitor w/Powered Speakers, \$1,004.00
- 2 9GB Seagate 3.5" Ultra2 LVD SCSI Disk Drive, \$1,112.00

4 Key Personnel Changes

- On October 1, Peter Lee started a temporary leave of absence from Carnegie Mellon University. During his leave, he will be working on a startup company to commercialize Proof-Carrying Code technology that was developed in the Fox Project. This startup company, called Cedilla Systems Incorporated, is funded in part by DARPA. He plans to return to CMU and the Fox Project in the Fall of 1999.

5 Noteworthy Meetings

- None

6 Administrative Data

Base Funding (excludes options): 5,630,798

Funded Options:

UNFunded Options: 648,704

Total Funding Provided to Date (both base and options): 3,617,974

Total Funding Expended to Date (both base and options): 3,164,416

Total Funding UNExpended: 453,558

Date Current Funding will be Expended: 30 APR 1999

Funding Expended in Most Recent Quarter: 247,490

Incremental Funding required for FY 1999: 800,000

Date of Financial Data: 31 DEC 1998